

BRIDGING NONLINEARITIES AND STOCHASTIC REGULARIZERS WITH GAUSSIAN ERROR LINEAR UNITS

Dan Hendrycks*
University of Chicago
dan@ttic.edu

Kevin Gimpel
Toyota Technological Institute at Chicago
kgimpel@ttic.edu

Abstract

We propose the Gaussian Error Linear Unit (GELU), a high-performing neural network activation function. The GELU nonlinearity is the expected transformation of a stochastic regularizer which randomly applies the identity or zero map, combining the intuitions of dropout and zoneout while respecting neuron values. This connection suggests a new probabilistic understanding of nonlinearities. We perform an empirical evaluation of the GELU nonlinearity against the ReLU and ELU activations and find performance improvements across all tasks.

1 Introduction

The sigmoid function was once the most popular nonlinear activation function for neural networks. Despite having a probabilistic interpretation, it has fallen out of favor due to slow and inaccurate convergence. In contrast, the widely-used ReLU activation usually demonstrates superior convergence yet is lacking in terms of its probabilistic interpretation [10]. Despite these differences, both nonlinearities can have a tremendous impact on the overall performance of a neural network, and both activations yield greater outputs for greater inputs. Their utility and difference is how they preserve or diminish inputs.

Some neural network regularization strategies, in contrast, act irrespectively of a neuron's value. For example, dropout randomly sets neuron values to zero, and zoneout randomly preserves previously computed neuron values [13, 7]. Dropout can act as a stochastic zero map, and zoneout can act as a stochastic identity map. The ReLU is a deterministic zero map or identity map, depending on the input value. But this determinism of the ReLU is limiting. The stochasticity of dropout and zoneout may allow for a pseudo-ensemble [1] of networks with stochastic width or stochastic depth, respectively, and this leads to marked performance improvements [5, 14]. Since ReLUs lack stochasticity and since the aforementioned regularizers are irrespective of their input, the innovations have remained distinct even though each uses zero or identity maps. In this paper, we bridge the gap between nonlinearities and stochastic regularizers by considering a new stochastic regularizer that is dependent upon input values. Afterward, we encapsulate the stochastic regularizer into a deterministic activation function that we call the Gaussian Error Linear Unit (GELU). In experiments across several tasks, we find GELU activations to outperform both ReLU and ELU activations.

*Work done while the author was at TTIC. Code available at github.com/hendrycks/GELUs




Figure 1: The Gaussian Error Linear Unit ($\mu = 0, \sigma = 1$), the Rectified Linear Unit, and the Exponential Linear Unit ($\alpha = 1$).

2 GELUs and the Stochastic 0- I Map

We begin by describing the stochastic regularizer that underlies the GELU nonlinearity. We call the regularizer the stochastic 0- I map (SOI map). Before expounding on this map, we define a few variables to ground the discussion. To start, assume that a layer of a neural network receives input data $X \in \mathbb{R}^{\text{batch size} \times n_{\text{in}}}$ such that $X_{ij} \sim \mathcal{N}(0, 1)$. Now, we have seen that both the sigmoid and ReLU preserve larger inputs and diminish smaller inputs. Meanwhile, ReLUs, dropout, and zoneout apply zero or identity maps. In an effort to combine these two properties, we define the SOI map as randomly applying the identity transformation to an input with probability $\Phi(X_{ij}) = P(Y \leq X_{ij})$, $Y \sim \mathcal{N}(0, 1)$ and applying the zero map otherwise. The implication is that, as x increases, the probability that we “drop” the input decreases. Such a regularizer retains nondeterminism but maintains dependency upon the input value.

We can extract a deterministic function from this regularizer to obtain a traditional nonlinearity. A SOI map’s expected transformation to an input x is $\Phi(x) \times Ix + (1 - \Phi(x)) \times 0x = x\Phi(x)$. Loosely, this expression states that we scale x by how much greater it is than other inputs. We now make an obvious extension. Since the cumulative distribution function of a Gaussian is computed with the error function, we define the Gaussian Error Linear Unit (GELU) as

$$\text{GELU}(x) = xP(X \leq x)$$

where $X \sim \mathcal{N}(\mu, \sigma^2)$. Both μ and σ are possibly parameters to optimize, but throughout this paper we simply let $\mu = 0$ and $\sigma = 1$.

The GELU bears semblance to the ReLU and ELU (see [3] for an ELU description). For example, as $\sigma \rightarrow 0$ and if $\mu = 0$, the GELU becomes a ReLU. More, the ReLU and GELU are equal asymptotically. Unlike the ReLU, the GELU and ELU can be both negative and positive. In fact, if we used the cumulative distribution function of the standard Cauchy distribution, then the ELU (when $\alpha = 1/\pi$ or $\text{ELU}(x) = \mathbb{1}(x > 0)x + \frac{1}{\pi}\mathbb{1}(x \leq 0)(e^x - 1)$, where $\mathbb{1}$ is the indicator function) is $xP(C \leq x)$, $C \sim \text{Cauchy}(0, 1)$ asymptotically. These are some fundamental relations to previous nonlinearities.

However, the GELU has several notable differences. This non-convex, non-monotonic function is not linear in the positive domain and exhibits curvature at all points. Meanwhile ReLUs and ELUs, which are convex and monotonic activations, are linear in the positive domain and thereby can lack curvature. As such, increased curvature and non-monotonicity may allow




Figure 2: MNIST Classification Results. Left are the loss curves without dropout, and right are curves with a dropout rate of 0.5. Each curve is the the median of five runs. Training set log losses are the darker, lower curves, and the fainter, upper curves are the validation set log loss curves.

GELUs to more easily approximate complicated functions than can ReLUs or ELUs. In addition and significantly, the GELU has a probabilistic interpretation given that it is the expected SOI map, which combines ideas from dropout and zoneout.

3 Experiments

We evaluate the GELU, ReLU, and ELU on MNIST classification (grayscale images with 10 classes, 60k training examples and 10k test examples), MNIST autoencoding, and CIFAR-10 classification (color images with 10 classes, 50k training examples and 10k test examples) tasks. We do not evaluate nonlinearities like the LReLU because of its similarity to ReLUs (see [8] for a description of LReLUs).

3.1 MNIST Classification

Let us verify that this nonlinearity competes with previous activation functions. To this end, we train a fully connected neural network with GELUs ($\mu = 0, \sigma = 1$), ReLUs, and ELUs ($\alpha = 1$). Each 7-layer, 128 neuron wide neural network is trained for 50 epochs with a batch size of 128. We use the Adam optimizer and its suggested learning rate of 0.001 [6]. The weights are uniformly initialized on the unit hypersphere, as this has positive impact on each nonlinearity’s performance [4, 9, 11]. Last, note that we perform this task with no dropout and a dropout rate of 0.5. Figure 2 shows that the GELU tends to have the lowest median training log loss under both dropout rates. Consequently, although the GELU is inspired by a different stochastic process, it comports well with dropout.

3.2 MNIST Autoencoder

We now move to a self-supervised setting and train a deep autoencoder on MNIST. To accomplish this, we use a network with layers of width 1000, 500, 250, 30, 250, 500, 1000, in that order. We again use the Adam optimizer and a batch size of 64. Our loss is the mean squared




Figure 3: MNIST Autoencoding Results. Each curve is the median of three runs. Left are loss curves for a learning rate of 10^{-3} , the middle figure is for a 10^{-4} learning rate, and right is for 10^{-5} .




Figure 4: CIFAR-10 Shallow CNN Results. Each curve is the median of three runs. Left is the network without dropout, and right is the network with dropout. Each training set log loss curve is selected from the best curve over the learning rates 10^{-3} , 10^{-4} , 10^{-5} .

loss. We vary the learning rate from 10^{-3} to 10^{-5} .¹ The results are shown in Figure 3. The GELU either ties or significantly outperforms the other nonlinearities. This shows the GELU nonlinearity is both stable and accurate under different learning rates.

3.3 CIFAR-10

Next, we demonstrate that for more intricate architectures the GELU nonlinearity again outperforms other nonlinearities. Using the CIFAR-10 dataset, we evaluate this activation function on shallow and deep convolutional neural networks. With a shallow architecture, our network has the stacks $(2 \times 3 \times 32)$, $(2 \times 3 \times 64)$ representing the number of layers, receptive field, and number of filters, respectively. We then feed this output through a two layer network with 512 and 256 neurons for the two layers. We apply max-pooling after every stack, and we run two

¹Afterward, we tried with a 10^{-2} learning rate but all solutions poorly converged, GELUs tied with ReLUs, and ELUs diverged.




Figure 5: CIFAR-10 VGG Net Results. Each training set log loss curve is selected from the best curve over the learning rates 10^{-3} , 10^{-4} , 10^{-5} .

experiments: we apply no dropout or use dropout rates of 0.25 after the first stack, 0.25 after the second stack, and 0.5 before the last fully-connected layer. As ever, we use the Adam optimizer, tune over the learning rates 10^{-3} , 10^{-4} , 10^{-5} , and initialize weights on the unit hypersphere (each filter has an ℓ_2 norm of one). Figure 4 shows the results. In both situations, GELUs provide faster and superior convergence.

Similar convergence gains can be found in a VGG Net-like [12] architecture. It has the stacks $(2 \times 3 \times 64)$, $(2 \times 3 \times 128)$, $(3 \times 3 \times 256)$, $(3 \times 3 \times 512)$, $(3 \times 3 \times 512)$ followed by two fully-connected layers, each with 512 neurons. To regularize the deep network, we drop out 40% of the neurons for every layer except the first two and last two layers, where we drop out no neurons and 50% of the neurons the first two and last two layers, respectively. Max pooling occurs after every stack, the learning rate was tuned over $\{10^{-3}, 10^{-4}, 10^{-5}\}$, and we decay the learning rate by 0.1 every 50 epochs all while training for 200 epochs with the Adam optimizer and unit hypersphere initialization. The results are shown in Figure 5. In summary, the GELU surpasses the ELU and ReLU for our shallow and deep convolutional neural networks.

4 Discussion

Across several experiments, the GELU outperformed previous nonlinearities, but there are minor cautions. The MNIST experiments are similar to those run by the architects of the ELU, but unlike their experiments, we use Adam instead of ordinary gradient descent because Adam better reflects what is typically used in practice for training neural networks. Interestingly, this change causes ELUs to occasionally perform worse than ReLUs. Furthermore, by the GELU’s difference in convexity and monotonicity, we find that an optimizer like Adam is important for allowing GELUs to converge well, while previous nonlinearities do not have this (minor) drawback. This suggests GELUs generate an error surface unlike those generated by ReLUs. Also, like $\exp(\cdot)$ and $\sigma(\cdot)$, the error function is expensive to compute accurately, but fast approximations exist. For example, $0.5x(1 + \tanh[\sqrt{2/\pi}(x + 0.044715x^3)])$ is a sufficiently fast, easy-to-implement approximation [2].²

It is also worth mentioning why we did not display the SOI map performance in the above

²Thank you to Dmytro Mishkin for bringing an approximation like this to our attention.




Figure 6: MNIST Classification Results with the SOI map evaluated. The dropout rate is 0.5, but when we apply the SOI map, we do not use any dropout. When feeding forward the validation set, we replace the SOI map with a GELU since the GELU is deterministic.

experiments. We now show one such example of its performance in Figure 6. Clearly, it can replace nonlinearities occasionally. For example, on the MNIST Classification tasks, it competes with the GELU and therefore outperforms other traditional nonlinearities—a neural network with only a regularizer can outperform a network with a nonlinearity. It achieves low validation log loss as a consequence of its regularizing stochasticity, and, in this case, regularizes the network better than dropout and a traditional nonlinearity. However, how best to adjust the regularization strength of the SOI map is currently unclear, so we cannot yet recommend using the SOI map as a nonlinearity replacement.

5 Conclusion

We observed that the GELU outperforms previous nonlinearities across several tasks. Importantly for future work, this nonlinearity has a probabilistic interpretation and may thus lead to a deeper understanding of the feedforward process as a whole. Other future research avenues include researching the LaLU, $xP(L \leq x)$, $L \sim \text{Laplace}(0, 1)$, as this may encourage neuron sparsity. Another avenue is finding a useful way to apply the SOI map across many different tasks because it remains unclear how best to modify the extent to which the SOI map regularizes. Fortunately, the GELU does not require any apparent adjustment to exceed the accuracy of previous nonlinearities.

Acknowledgment

We would like to thank the NVIDIA Corporation for donating GPUs used in this research.

References

- [1] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *Advances in Neural Information Processing Systems*, pages 3365-3373, 2014.

- [2] Amit Choudhury. A Simple Approximation to the Area Under Standard Normal Curve. In *Mathematics and Statistics*, 2014.
- [3] Djork-Arne Clevert, Thomas Unterthiner, Sepp Hochreiter. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). In *International Conference on Learning Representations*, 2016.
- [4] Dan Hendrycks, Kevin Gimpel. Generalizing and Improving Weight Initialization. In arXiv.
- [5] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, Kilian Weinberger. Deep Networks with Stochastic Depth. In arXiv:1603.09382.
- [6] Diederik P. Kingma and Jimmy Lei Ba.. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*, 2015.
- [7] David Krueger, Tegan Maharaj, Jnos Kramr, Mohammad Pezeshki, Nicolas Ballas, Nan Rosemary Ke1, Anirudh Goyal, Yoshua Bengio, Hugo Larochelle, Aaron Courville, and Chris Pal. Zoneout: Regularizing RNNs by Randomly Preserving Hidden Activations. In arXiv:1606.01305.
- [8] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *International Conference on Machine Learning*, 2013.
- [9] Dmytro Mishkin and Jiri Matas. All You Need Is a Good Init. In *International Conference on Learning Representations*, 2016.
- [10] Vinod Nair and Geoffrey E. Hinton. Rectified Linear Units Improve Restricted Boltzmann Machines. In *International Conference on Machine Learning*, 2010.
- [11] Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *International Conference on Learning Representations*, 2014.
- [12] Karen Simonyan, Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations*, 2015.
- [13] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. In *Journal of Machine Learning Research*, 2014.
- [14] Andreas Veit, Michael Wilber, and Serge Belongie. Residual Networks are Exponential Ensembles of Relatively Shallow Networks. In arXiv:1605.06431.